

 Open access • Posted Content • DOI:10.1101/2020.07.27.223057

The ABCD Stop Signal Data: Response to Bissett et al. — [Source link](#)

Hugh Garavan, Sage Hahn, Bader Charani, Anthony C. Juliano ...+9 more authors

Institutions: University of Vermont, University of Michigan, Swinburne University of Technology, Yale University ...+2 more institutions

Published on: 28 Jul 2020 - bioRxiv (Cold Spring Harbor Laboratory)

Related papers:

- [The Adolescent Brain Cognitive Development \(ABCD\) study: Imaging acquisition across 21 sites.](#)
- [fMRI visualization of brain activity during a monetary incentive delay task.](#)
- [Function in the human connectome: Task-fMRI and individual differences in behavior](#)
- [Balancing precision with inclusivity in meta-analyses: A response to Roos and colleagues \(2017\)](#)
- [Where is the criterion noise in recognition? \(Almost\) everywhere you look: comment on Kellen, Klauer, and Singmann \(2012\).](#)

Share this paper:    

View more about this paper here: <https://typeset.io/papers/the-abcd-stop-signal-data-response-to-bissett-et-al-540z8sryo6>

The ABCD Stop Signal Data: Response to Bissett et al.

Garavan H^{1*}, Hahn S¹, Chaarani B¹, Juliano A¹, Allgaier N¹, Yuan DK¹, Weigard A², Orr C³, Watts R⁴, Wager TD⁵, Ruiz de Leon O⁶, Hagler DJ, Jr⁶, Potter A¹.

¹Department of Psychiatry, University of Vermont, Burlington VT, USA

²Department of Psychiatry, University of Michigan, Ann Arbor MI, USA

³Department of Psychology, Swinburne University, Melbourne, Australia

⁴School of Medicine, Yale University, New Haven CT, USA

⁵Department of Psychological and Brain Sciences, Dartmouth College, Hanover NH, USA

⁶Department of Radiology, University of California, San Diego, La Jolla, CA, USA

* Corresponding Author:

Department of Psychiatry,
University of Vermont, UHC Building,
1 South Prospect Street,
Burlington VT,
05401
USA

E-mail: hgaravan@uvm.edu

Acknowledgements

We thank Frederick Verbruggen and members of the ABCD task fMRI working group for their helpful discussions of these topics.

Abstract

This paper responds to a recent critique by Bissett and colleagues (<https://www.biorxiv.org/content/10.1101/2020.05.08.084707v1>) of the fMRI Stop task being used in the Adolescent Brain Cognitive Development (ABCD) study. The critique focuses primarily on a design feature of the task that the authors contend might lead to a violation of race model assumptions (i.e., that the Go and Stop processes are fully independent) which are relevant to the calculation of the Stop Signal Reaction Time, a measure of the inhibition process. Bissett and colleagues also raise a number of secondary concerns. In this response we note that satisfying race model assumptions is a pernicious challenge for Stop task designs but also that the race model is quite robust against violations of its assumptions. Most importantly, while Bissett et al. raise conceptual concerns with the task we focus here on analyses of both the performance and the neuroimaging data and we conclude that the concerns appear to have minimal impact on the task data. We note that there were errors in the analyses conducted by Bissett et al. that inflated their estimates of race model violations in ABCD, that they did not apply any performance-based exclusions to the data they analyzed, that a number of the errors that they flagged were already identified and corrected in the ABCD annual data releases, and we argue that a number of the other concerns reflect sensible design decisions. In this paper, we list some adjustments that will be made to the task and some new flags that will be added to the annual, curated data releases. We stress that the ABCD data are fully available to the scientific community who are empowered to apply whatever inclusion and exclusion criteria they deem appropriate for their analyses and we conclude that the ABCD Stop task yields valuable data that researchers can use to track adolescent neurodevelopment.

Bissett and colleagues (Bissett, Hagan et al., Under Review) list a number of concerns regarding the specific version of the Stop Signal Task that is included as one of three fMRI tasks in the neuroimaging battery of the multi-site, longitudinal Adolescent Brain Cognitive Development (ABCD) study (www.ABCDstudy.org). The ABCD study is committed to full sharing of its tasks, data processing and analysis scripts, and dataset. We are delighted to see “open science” in action, and we value the scientific community flagging concerns and helping us to continually improve this landmark study.

The Bissett et al. critique focuses largely on stimulus design characteristics of the ABCD Stop Signal Task (SST) which may lead to violations of race model assumptions (Logan et al., 1984). The concerns are well described in Bissett et al. In brief, the onset of the Stop signal, which is controlled by an adaptive, performance-related algorithm, is accompanied by the simultaneous offset of the Go choice stimulus. One consequence of this design feature is that the on-screen duration of the Go choice stimulus can be shorter on Stop trials. The concern raised by Bissett et al. is whether this might affect the context independence of the Go process, that is, that the Go process is unaffected by whether or not a Stop signal is presented on a trial. Context independence is assumed by race model theories of response inhibition and underpins the valid calculation of the Stop Signal Reaction Time (SSRT), an estimate of the duration of the stopping process. In addition to this primary issue, Bissett and co-authors raise a number of other concerns with the task design and conclude that these “significantly limit” the value of the data. We disagree with this conclusion and provide evidence in support of the utility and validity of the SST data from the ABCD Study. The Bissett et al. critique raises important theoretical issues related to the assumptions of the race model underlying calculation of the SSRT, but the issues raised do not necessarily undermine the utility of the ABCD SSRT estimates as a measure of individual differences. Importantly, Bissett et al. do not examine whether the issues they raised do in fact affect the ABCD study’s SSRT estimates, nor whether they have an impact on the brain imaging data.

While acknowledging the concerns regarding race model violations, we focus here on empirically investigating the extent to which these violations meaningfully impact the quality of the ABCD data. We present a series of analyses of both the SSRT estimates and the validity of the neuroimaging data. Ensuring independence of the Go and Stop processes is a perennial concern with the SST and not one peculiar to the ABCD task version (Bissett, Poldrack et al., Under Review). Moreover, the SST is quite robust against violations of race model assumptions (Band et al., 2003). Thus, it is critical to determine whether, or to what degree, the concerns raised do indeed corrupt the SSRT estimates. We demonstrate here that they have only modest effects on SSRT and, furthermore, do not substantively impact or invalidate the SST fMRI measures.

Some of the concerns raised by Bissett et al. were in fact known issues that were corrected prior to the annual data analyses and every curated data release of the ABCD Data Analysis, Informatics and Resources Center. We note also an error in the analyses conducted by Bissett et al. that led to a substantial overstatement of the extent of race model violations within the ABCD data. Bissett et al. applied no performance flags (either their own or those recommended in all ABCD data releases) that serve to exclude participants who do not perform the task appropriately. Other recommended corrections, such as flipping left-right responses in those participants with reversed response paddles, were not applied. And, finally, we contend that many of the minor issues raised by Bissett et al. are design features and not design flaws. Following the format of the original paper, we address each issue in turn and offer our recommendations on if, and how, each might be addressed. The ABCD task fMRI working group is currently deciding on changes that might be made to the task taking into consideration the impact of the issues raised by Bissett

et al. and balancing these against the implications of changing a task once a longitudinal study has commenced.

Sample selection and basic SSRT characteristics. Before addressing the issues raised by Bissett et al., we first describe some aspects of the sample and the SST data. In order to analyze the same subset of ABCD participants analyzed by Bissett et al., we downloaded the data of 8,809 of their initial 8,811 participants from the FastTrack Recommended Active Series. 8,450 participants had both valid SST data and performance summary scores in the 2.0 Data Release. This number does not match Bissett et al.'s final number (7,231) due, we believe, to Bissett et al. missing available .csv files in their download (see "Match_Data.ipynb" on our github page). We do not believe this is consequential as performance summary statistics across the two subsets of participants are very similar. The code underlying all analyses conducted in this paper are available (https://github.com/sahahn/SST_Response). We underscore that although the annual ABCD data releases provide updates and contain flags identifying problematic data (e.g., poor quality images, poor task performance), the data are fully available to the scientific community, empowering researchers to apply whatever inclusion and exclusion criteria they deem appropriate for their specific research questions. Figure 1 shows the performance criteria that have been used to calculate the SST performance flag that is part of the recommended inclusion criteria for these data provided with each ABCD data release.

Figure 1 shows two important properties of the ABCD SST data. First, performance statistics show that the tracking algorithm achieves a successful inhibition rate of approximately 50% (which maximizes efficiency of SSRT estimation), few omission errors, and typically distributed response times. Second, it shows robust stopping-related activation in known response inhibition-related regions. These characteristics increase confidence that the ABCD SST data offer insightful measures of inhibitory control abilities and related brain function.

Performance Flags	
Correct Go	< 60%
Incorrect Go	> 30%
Late Go	> 30%
Go Omissions	> 30%
Total no. Go trials	< 300
Stop Success Rate	< 20% or > 80%

Group Performance	
Stop Success Rate	51.5%
Go Omission Rate	4.6%
Go RT	544 msec
Stop Fail RT	460 msec
Go Accuracy	91.0%
Stop Fail Accuracy	80.8%
SSRT	283±78


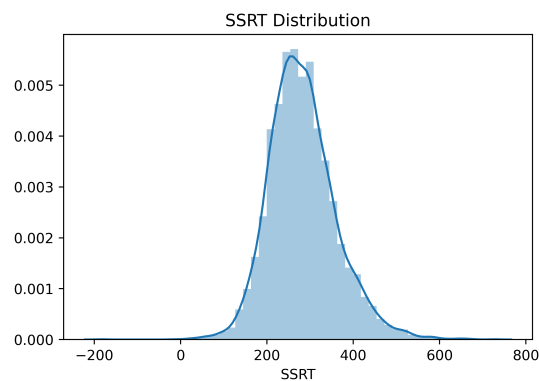



Figure 1: ABCD Stop Signal Task. The performance criteria (recommendations for participant exclusion) and group level performance on the ABCD Stop Signal Task are shown in the tables. Histograms of the p(Successful Stopping) and SSRT and group activation maps (Cohen's d threshold of .2 for Successful Stops vs Correct Go trials), for the baseline (age 9 and 10) data are shown.

Issue 1. Different go stimulus duration across trials. As noted above and described in Bissett et al., the offset of the Go choice stimulus coincident with the onset of the Stop signal may reduce the strength of the Go process on Stop trials, violating context independence. Poorer choice accuracy on Stop Fail trials compared to Go trials (Figure 1), especially at shorter Stop Signal Delays, indicates that this may be a valid concern; the poorer choice accuracy on the Stop Fail trials is consistent with the Go process being different on these trials in comparison to Go trials. Bissett et al. report accuracy of 89% and 78% for Go and Stop Fail trials, respectively. In our analysis, with performance criteria applied and the exclusion of a small number of participants

who experienced a task programming error (see Issue 3 below) Go accuracy = 91% and Stop Fail accuracy = 81%.¹

The critical issues are to what extent this violation of race model assumptions impacts reaction time and/or brain imaging data, and whether any violation impacts the utility of SSRT as a measure of individual differences in inhibitory control. Full context independence is difficult to attain. Bissett et al. and others demonstrate that this is the case across many Stop Signal Tasks, including those without the stimulus design feature of the ABCD task (Gulberti et al., 2014; Bissett, Poldrack et al., Under Review). The presentation of a Stop signal has the potential to impact an ongoing Go response, even if the Go stimulus remains on screen. In addition, it is known that participants slow their Go responses in anticipation of a Stop trial (indeed, this “proactive” control can be modelled; Harlé et al., 2016). “Strategic” adjustments in the speed of Go trial responding can vary within an experiment and across individuals, and can affect SSRT estimates (Leotti et al., 2010).

One test for clear violations of context independence, which Bissett et al. presents, is whether Stop Fail response times (RTs) are slower than Go RTs. This pattern cannot be explained by the standard race model, which assumes Go processes are equivalent on Go and Stop signal trials and, further, that the observed distribution of RTs on Stop Fail trials are censored (by virtue of the Stop process completing before the relative slower Go responses would be made). As Bissett et al. note, the ABCD data pass this test: Stop Fail trials are not slower than Go response times. Table 1A shows RT data for ABCD as calculated by Bissett et al., after we fixed errors in the Bissett et al. code², applied the standard performance flags, and excluded participants affected by the Issue 3 programming error (described below). Stop Fail RT is 83 msec faster than Go RT ($t_{(1,7114)} = 53$, $p < .0001$), broadly consistent with context independence. For comparison, we include two datasets from Bissett and colleagues, chosen because they do not share the ABCD design property of offset of the Go stimulus with the onset of the Stop stimulus. We refer to these datasets as the Ontology study ($n = 522$; Eisenberg et al., 2019) and the Phenome study ($n = 130$, healthy controls only; Poldrack et al., 2016). We use the low frequency condition (20% Stop trials) in the Ontology study as it is closer to the ABCD proportion of 17%. The ABCD value of 83 msec is within the range of these other studies: Stop Fail is 30 msec and 122 msec faster than Go trials in the Ontology and Phenome studies, respectively.

¹ Excluding participants who fail to adequately comply with task instructions or who show aberrant performance is a standard practice and especially important given the young age of ABCD participants. Recommended performance criteria that accompany ABCD data releases are shown in Figure 1 but these were not applied by Bissett et al. In addition, a programming error detailed under Issue 3 led to a very different task experience for certain participants and we are recommending that these participants (1.24% of the sample) be excluded from analyses.

² During Stop trials, responses can occur during the Stop Signal Delay (SSD; that is, before the Stop signal is presented), during presentation of the Stop signal, or during the subsequent fixation period. A response recorded in any of these intervals results in a Stop Fail trial. In calculating an adjusted Stop Fail response time, Bissett et al. correctly added the duration of the SSD to responses made either during the Stop signal or the subsequent fixation period. However, they also added the full SSD duration to all responses made during the SSD itself. This error results in inflated Stop Fail response times for these trials.

ABCD					Comparison	
(A)	Initial	Stop Fail RT corrected	Exclude poor performers	Issue 3 exclusion	Ontology	Phenome
Mean Go RT	543	543	543	544	571	478
Mean SF RT	495	459	458	460	541	356
Difference	48	84	85	83	30	122

(B)	Initial	Stop Fail RT corrected	Exclude poor performers	Issue 3 exclusion	Ontology	Phenome
%Violators	18.4	7.1	4.9	5	17	4.4

(C)	ABCD	Ontology	Phenome
Significant Violators	1.6%	2.7%	0%

Table 1: Performance statistics relevant to Issue 1 for ABCD and two comparison datasets (see text for details). (A) We replicate the initial calculations by Bissett and colleagues of the mean response times for Go RT and Stop Fail (SF) RT trials, we correct their calculation error for SF RT, we apply the exclusion of participants based on poor performance as recommended with the ABCD annual data releases, and then we drop the RT data of an additional 1.24% of participants whose task was compromised by a coding error (described under Issue 3 below). (B) We identify the percentage of participants who do not show the expected pattern of SF RT < Go RT following the same steps as in A and we calculate the percentage of participants who show significant differences between SF RT and Go RT. As can be seen in the final row the percentage of violators in the ABCD study is in line with that of other studies as is quite small, 1.6%.

Next, Bissett et al. reports that 18% of the ABCD participants do not show the expected RT pattern and label them as violators (i.e., participants for whom Stop Fail RT > Go RT). Table 1B shows that the correct estimate is 5% when the appropriate corrections are applied, and this is superior or comparable to the estimates for the Ontology and Phenome studies (17% and 4.4%, respectively). Furthermore, these percentages may not indicate statistically reliable effects (i.e., the differences in RTs may not differ reliably from 0), as confidence intervals and estimates of their sampling distribution are not provided. Although analyses of individual participants are likely underpowered, with too few trials for robust behavioral analyses, Table 1C reports the numbers of participants who might be deemed true violators (i.e., with a significant one-tailed t-test per participant comparing Stop Fail trial RTs against Go trial RTs). The percentage is low: 1.6% for ABCD, again falling between the Ontology and Phenome studies that were designed without coincident Go offset and Stop onset.

We conclude that some number of participants evidencing context independence violations are to be expected in many SST designs. While researchers can make their own decisions on participant inclusion and exclusion criteria given ABCD's data sharing procedures, we note that a recent SST "best practices" paper recommends that SSRT should not be estimated for those participants who violate the Stop Fail RT < Go RT criterion (Verbruggen et al, 2019).

Consequently, a new flag identifying the 5% of ABCD participant “violators” is included in the ABCD 3.0 release (see Implications and Recommendations below). Note that this flag is applied to any participant whose Stop Fall RT > Go RT by any amount (i.e., 1 ms or greater).

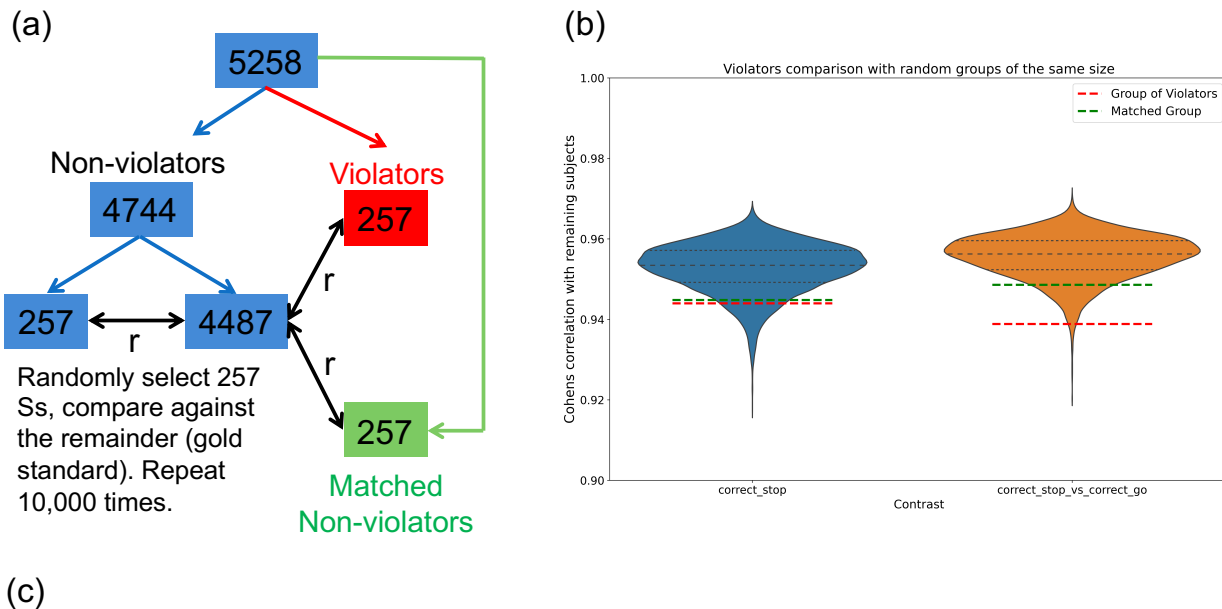
Examining individual differences is a central goal of the ABCD study. Crucially, for the utility of the ABCD SSRT measure to be degraded as a measure of individual differences, violations of context independence must result in more than a shift in mean SSRT. Rather, the rank ordering of participants’ SSRT values must be substantively altered. There is no evidence provided that this is the case and, indeed, as all ABCD participants performed the same task, we expect individual differences to be largely unaffected by the abbreviation of the Go process as a function of task structure. Future experimental or computational modeling studies could help provide greater clarity about these matters. Until that time, researchers are encouraged to carefully consider the assumptions of any measurement model they apply to these data, including the race model-based SSRT estimate, and to consider the possible limitations of parameter estimates and measures derived from any model that assumes context independence.

Turning to the brain activation data, it is important to note that the measurement assumptions that the activation contrasts reflect valid measures of response inhibition are much simpler than those required for the SSRT estimation and do not rely on race model assumptions, including context independence. Brain activity can be associated with response inhibition processes if one compares trials requiring inhibition of prepotent responses against trials that do not. A standard contrast to achieve this end for ABCD would be to compare Successful Stop trials against Go trials. The shorter duration of Go choice stimuli on Stop trials compared to Go trials does introduce differences between the two conditions. However, there are typically other more substantial differences present when isolating inhibition-related activation: There is a motor response on Go trials and not on Stop success trials, only the latter contains a Stop signal, and so on. The contrast of Successful Stop trials (against the implicit baseline) is also available to researchers. As is always the case, researchers using these data should be aware of design specifics and determine if they impact on the researcher’s specific question. The ABCD SST has already been shown to produce robust activation in the response inhibition network and activation levels show the anticipated correlations with individual differences in SSRT (Casey et al., 2018; Chaarani et al., Under Review). A very similar task, with the same Go stimulus design features, has been employed in the IMAGEN study of adolescent development (Schumann et al., 2010) and has, for example, identified functional differences between adolescents with substance use, adolescents with ADHD, adolescents with psychotic symptoms, dysregulated youth and controls (Bourque et al., 2017; Specher et al., 2019a; Whelan et al., 2012), and has predicted future drug use (Spechler et al., 2019b; Whelan et al., 2014).

We demonstrate the validity of the brain activation measures in ABCD with two analyses. The first examined brain activation in the violators identified above (a “worst case” scenario in which we might expect atypical activation patterns) and the second examined the impact of removing trials with short Stop Signal Delays (SSD) from all participants. Short SSD trials (SSD < 150 msec), in which the Go stimuli were presented for the shortest durations, are likely to be the trials driving any context independence violations. The first analysis compared Stop activation (the contrasts Successful Stops and Successful Stops vs Correct Go trials) between the group of violators ($n = 257$; baseline [ages 9 and 10] data only) and non-violators ($n = 5,001$). Covariates included sex, age (in months), highest parental education, race/ethnicity, puberty level, and scanner. No group differences were observed with a vertex-wise threshold of $p < .05$ and application of either a family-wise error correction based on random field theory or a less conservative false discovery rate. At a nominal $p < .05$, uncorrected threshold, group differences were observed in visual

cortex only. Further, in a region-of-interest analysis of the right IFG, a critical node of the response inhibition network, violators and non-violators did not differ in activation ($p = .83$).

To quantify similarity between groups in the spatial patterns of activation, we calculated the vertex-wise correlation between group activation of the violators and the “gold standard” of activation based on the remaining non-violator participants (the whole sample was first residualized for the covariates listed above). To facilitate interpretation, we quantified the similarities that would be expected with samples of size 257 by comparing randomly selected subsamples of non-violators ($n = 257$, 10,000 samples) against the remainder ($n = 4,487$; Figure 2A). Figure 2b shows the distribution of vertex-wise correlations for both the Successful Stops contrast (blue) and the Successful Stops vs Correct Go contrast (orange). The correlations for violators are indicated by the red line. Although lower than the mean of the subsampling distribution, shown by the dotted blue line, we note that the vertex-wise correlation is very high, even for violators ($r = .94$ and $.94$ for violators, compared to the full group of non-violators with $r = .95$ and $.96$, for Successful Stops and Successful Stops vs Correct Go, respectively). Moreover, violators, unsurprisingly, show relatively poor performance on the task (Figure 2c). An equal sized group of non-violators, matched to the violators on Stop success rate, Stop Fail accuracy, Go RT, Go accuracy, and Go omission rate are indicated by the green lines (Matched Group). We conclude that even those participants identified as violating race model assumptions show activation patterns that are very similar to those observed for non-violators.



	Stop Success Rate	Stop Fail Accuracy	Go RT	Stop Fail RT	Go Accuracy	Go Omission Rate	Mean SSD
All Non-violators	52%	82%	547	454	92%	4%	236
Violators	46%	71%	533	614	89%	6%	139
Matched Non-violators	46%	71%	525	438	88%	6%	148

Figure 2: Correlation of violators and performance-matched non-violators with gold-standard brain activation. (a) A subsampling procedure determined the similarity between vertex-wise activation levels in samples of $n=257$. (b) Correlations with the “gold standard” activation map for subsamples of non-violators (blue and orange distributions), violators (red line) and performance-matched non-violators (green line). (c) Performance of violators, all non-violators, and performance-matched non-violators.

The second analysis compared group Stop activation maps (Successful Stops vs Correct Go) with all trials versus with the shorter SSD Stop trials (0msec, 50msec, 100msec) excluded. The same covariates as described above were included ($n = 5,058$). Although the amplitude of activation was larger in the former (due, presumably, to the inclusion of more trials), critically, the patterns of activation were almost identical (Figure 3). The vertex-wise correlation in group activation maps for violators versus non-violators was $r = .99$.

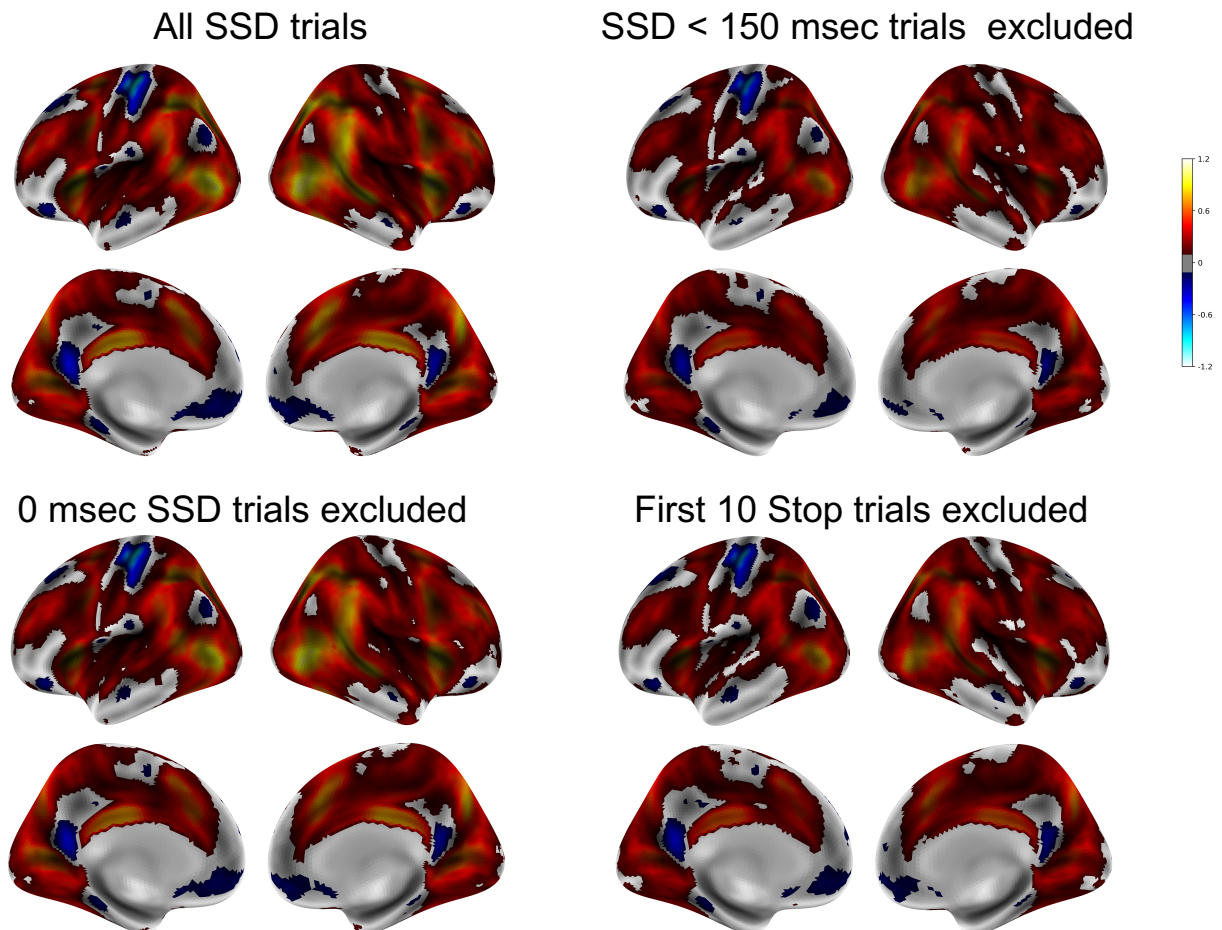


Figure 3: Brain activation (Cohen's d threshold of .2 for Successful Stops vs Correct Go trials) for All Stop trials, Stop trials with SSD = 0/50/100 msec excluded, Stop trials with SSD = 0 msec excluded, and first ten Stop trials excluded.

Issue 1 Implications and Recommendations: Race model violators (Stop Fail RT > Go RT) are identified in Release 3.0 and we suggest that researchers not include them when estimating SSRT (Verbruggen et al. 2019). We do not believe that there is sufficient evidence at this stage to warrant distrust of the remaining performance and neuroimaging data but encourage investigators to consider the impact of the issues that have been raised for their research question and their application of SSRT measurement models, and to apply what they deem to be appropriate inclusion and exclusion criteria. The ABCD task fMRI working group is currently considering whether the SST design should be changed. The poorer choice accuracy on Stop Fail trials indicates a degree of independence violation that likely will be reduced if the Go stimuli

remain on screen (for the 1 sec duration of the trial or until a response is made). However, the analyses reported above suggest that the current task is yielding valuable, useful data that would appear to be largely unaffected by the concerns raised by Bissett et al. Clearly, an ongoing longitudinal study will prioritize not changing a task without compelling, and ideally empirical, reasons to do so. If changes to the task are to be made then a new task design will be carefully phased in enabling us to assess any differences between designs.

Issue 2. Go stimulus sometimes not presented. Arising from the stimulus design feature underlying Issue 1, Go choice stimuli are not presented on trials where the SSD drops to 0 msec. Bissett et al. suggest that this may confuse participants or “may make successfully stopping trivial (as the go process never started).” Bissett et al. report that 9.1% of Stop trials are 0 msec SSD trials but once the performance flagged participants and the Issue 3 programming error participants are excluded this drops to 6.9%. Curiously, performance on these trials is not trivial (average successful inhibition rate across participants on these trials is 64.8%) reflecting, presumably, the response prepotency induced by the high proportion of Go trials (see Table 2) which suggests that these trials also engaged inhibitory control processes.

	All Ss	Exclude poor performers	Issue 3 exclusion
% of trials	9.1%	7.3%	6.9%
Stop Success Rate	60.3%	63.7%	64.8%

Table 2: The percentage of Stop trials in which the SSD = 0msec and the probability of successfully inhibiting a response on these trials.

The 0 msec SSD trials are broadly distributed across participants, with 49.9% of participants having at least one 0 msec SSD trials. We calculated SSRT with these 0 msec SSD trials included (283 ± 78) and excluded (276 ± 73). (The average of a participant's SSDs are included in the calculation of their SSRT, so the exclusion of all 0 msec SSD trials would be expected to produce a shorter SSRT estimate.) Notably, the correlation between the two estimates is very high ($r = .97$). The broad distribution of these trials across participants appears to reduce their impact on subsequent analyses. Turning to the neuroimaging data, the vertex-wise correlation between brain activation when these trials are included vs. excluded is very high ($r = .99$; see Figure 3). This correlation holds when all participants are included ($n = 5,064$) and when the analyses are restricted to those participants with one or more 0 msec SSD Stop trials ($n = 2,416$).

Issue 2 Implications and Recommendations: The impact of these trials on the existing data appears to be very small. For researchers who may wish to exclude participants with a high number of these trials, we include the number of 0 msec SSD trials per participant in Release 3.0. If the changes suggested under issue 1 are made to the task design then this will remove this Issue 2 concern going forward.

Issue 3. Degenerate stop-signal delays. Bissett et al. identified a programming error and we thank them for bringing this to our attention. When the SSD is 50 msec, a response that is faster than 50 msec is erroneously recorded as the response for all subsequent Stop trials. Bissett et al. report that this programming error affects 2.67% of participants. However, if the performance flags are applied this reduces to 1.24% of participants. The data of many of these participants are likely retrievable if one restricts analyses to the data obtained prior to the onset of the error. For example, just 0.8% of participants have this problem occur prior to their 50th Stop trial and the

“best practices” paper by Verbruggen and colleagues concludes from a series of simulations that “reliable and unbiased SSRT group-level estimates can be obtained with 50 stop trials.”

Issue 3 Implications and Recommendations: The programming error in the task is being corrected for future data collection and the corrected task will be available on the ABCD github (<https://github.com/ABCD-STUDY/>) and study website (www.ABCDStudy.org/Scientists/). For the existing data, although we anticipate that valuable performance and brain imaging data of many participants affected by this programming error are retrievable, we include a variable identifying these participants in Release 3.0, enabling researchers to exclude them from analyses.

Issue 4. Different Stop Signal duration for different SSDs. This issue arose because all trial events were constrained to happen within the 1 second trial period and not to carry into the inter-trial interval. As a consequence, the duration of the Stop signal (typically 300 msec) was shortened if the SSD was greater than 700 msec. One might expect these events to be quite uncommon and also to be indicative of other performance-related problems (i.e., a 700 msec SSD is atypically long). While Bissett et al. report that shorter Stop signal durations occur on 0.8% of Stop trials, we calculate the frequency to be 1.15% of all Stop trials. However, this reduces to 0.12% of trials once performance criteria are applied (as mentioned, the presence of these very long SSD trials indicate other performance problems). Moreover, as participants very often respond during the relatively long SSD, the proportion of trials in which the shorter SSDs are presented and on which participants have not already responded reduces to 0.07% of trials. We removed these short Stop signal duration trials and assessed the impact on SSRT. Identical results were obtained: With all trials included and with these short Stop signal duration trials included, $SSRT = 283 \pm 78$.

Issue 4 Implications and Recommendations: Given the rarity of these trials, we conclude that they have a negligible impact on the SSRT estimates. Nonetheless, we will change the task to ensure that the Stop signal duration is always 300 msec in duration as this change will have a negligible impact on any longitudinal comparisons.

Issue 5. Non-uniform conditional trial probabilities. Bissett et al. note correctly that the trial orders (Stop vs. Go) were not fully randomized (in accordance with the ratio of Stop to Go trials). Instead, the conditional trial probabilities and inter-stimulus intervals (ISIs) for ABCD were selected to optimize the joint estimation efficiency for fMRI responses to Go and Stop events, assuming a canonical hemodynamic response function. For each of 40 runs, sequences of Stop and Go stimuli, and randomized jittered intervals between them, were selected by generating 100,000 random trial sequences and ISIs (constrained to 0.7 – 2 sec, starting mean 0.9 sec), and choosing the design with the minimum mean Variance Inflation Factor. The design was constrained to avoid repeating successive Stop trials. The rationale for this is based on evidence that repeated Stops are easier, thereby reducing power and homogeneity in the demand on inhibition across trials (Bissett et al., 2012).

Bissett et al. raise concerns that the non-uniform trial probabilities might be learned by participants. Their analysis of post-Stop RTs indicate that this is not the case (i.e., initial post-Stop slowing does not transition, as the task progresses, to post-Stop speeding up) but they caution that it may be learned as the participants age. A very similar SST (i.e., no repeated Stop trials) has been used in the IMAGEN longitudinal study of ~2,000 adolescents. We assessed post-Stop behavior similar to Bissett et al. (post-Stop RT minus pre-Stop RT, separating all Stop trials into quartiles) when these participants performed the task at age 19, which was the second administration of the task to them. Figure 4 shows a very similar pattern in IMAGEN to what Bissett et al. report for ABCD. There is no evidence that participants at age 19 learn to speed up on post-Stop trials and, instead, post-Stop slowing is maintained.

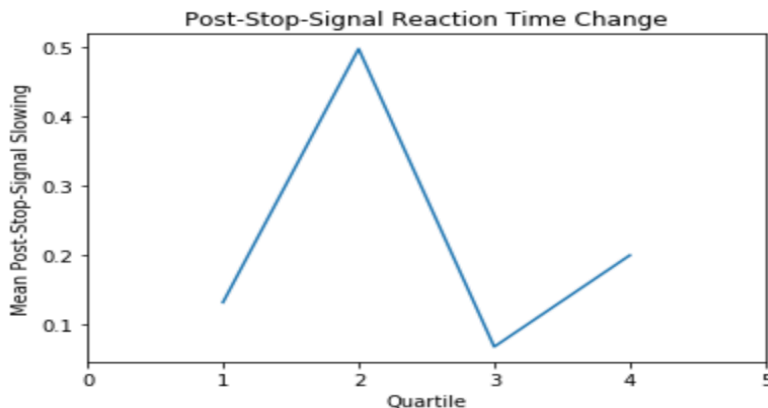


Figure 4: Post-Stop slowing (i.e., RT in seconds on the Go trial that immediately followed a Stop trial minus RT on the Go trial that immediately preceded the Stop trial) is shown across trial quartiles for the IMAGEN participants at age 19. This task also excludes repeated Stop trials and, similar to the younger ABCD participants, shows no evidence of changes in post-Stop slowing between the start and end of the task.

Issue 5 Implications and Recommendations: We recommend no changes to the task design. Researchers concerned about trial conditional probabilities being learned as the ABCD participants age can assess these in a manner as suggested by Bissett et al. Indeed, the extent to which participants develop subjective expectations and prepare for Stop trials is an interesting avenue of research for investigators interested in proactive control.

Issue 6. Trial accuracy incorrectly coded. Bissett et al. identify a number of trial outcome labeling errors. All of these errors were identified by the ABCD Data Analysis and Informatics Resource Center and were corrected prior to data analysis and data release. These errors derived from how E-Prime pre-release settings were used: Responses occurring during the pre-release period of the post-trial jittered intervals were not recorded correctly, resulting in Stop Fail trials being logged as Stop Successes and Correct Go trials being logged as Go Omissions. Details on the trial labeling errors and their correction are included in the data processing scripts available from the ABCD github site (<https://github.com/ABCD-STUDY/>). The misclassification of errors on the Stop trials (classifying what should be Stop Fails as Stop successes) did impact on the task as the subjects performed it (i.e., SSD increased when it should have decreased), but this occurred on only a small fraction of trials and, as Bissett et al. report, would have had a very small effect on the SSD tracking algorithm.

Issue 6 Implications and Recommendations: The cause of the errors has been identified and is being corrected in the task code. Data available through the annual NDA releases have already corrected for these labelling errors but researchers wishing to work with the raw E-Prime output are encouraged to employ the corrections specified in the “abcd_extract_eprime_sst” script available on the ABCD github site.

Issue 7. SSD values start too short. Bissett et al. query the reasoning behind starting the task with the SSD set to 50 msec rather than a value that is closer to what the final SSD would be (e.g., 250msec). The shorter SSD starts the task at a relatively easy level thereby providing a “warm-up” of sorts, easing participants, aged just 9 and 10 at baseline, into what is a cognitively challenging task. As shown in Figure 6 of Bissett et al., these first few trials are, in fact, not trivial

for participants insofar as performance is better than average but far from ceiling (starting at 75% and dropping to 60% by the fifth Stop trial). Consequently, these first few trials contribute usefully to the behavioral and activation measures. As shown in Figure 1, the 60 Stop trials of the ABCD task are sufficient for the adaptive algorithm to converge on ~50% Stop success rate. To assess the impact of these starting trials, we removed the first ten Stop trials and all Go trials up to the tenth Stop trial and assessed the impact on SSRT. With all trials included SSRT = 283 ± 78 and with the first ten Stop trials excluded SSRT = 284 ± 85 with a high correlation between the two ($r = .98$). Similarly, group brain activation with and without these first ten Stop trials was calculated ($n = 5,067$) with the vertex-wise correlation between the two being very high, $r = .99$ (see Figure 3).

Issue 7 Implications and Recommendations: We believe there are no implications for the task or the data arising from this design feature.

Issue 8. Low stop trial probability. The ratio of Go to Stop trials is 300:60 (17%) in ABCD, a ratio that increases task demands (i.e., increases the prepotency to respond, making inhibitions more difficult) which should ensure that this task continues to challenge participants as they age. We note that the task contains 60 Stop trials (as mentioned above, 50 are deemed sufficient for estimating SSRT) and successfully converges on ~50% Stop success rate (Figure 1). As noted above, a very similar task, with the same Stop trial probability has been used in the largest adolescent neuroimaging longitudinal study to date (IMAGEN; $n = 2,000$, with assessments at age 14, 19, and 23; Schumann et al., 2010) and it has demonstrated robust Stop success and Stop fail activations at ages 14 and 19 (age 23 data not yet available) and has demonstrated its ability to discriminate among adolescent phenotypes and predict future adolescent behavior (Whelan et al., 2012; 2014).

Issue 8 Implications and Recommendations: We believe there are no implications for the task or the data arising from this design feature. If researchers wish to compare the task's activation levels or SSRT estimates against other datasets then, as is always the case in comparing separate studies, they should be aware of the particular design features of the ABCD study.

Conclusion We encourage researchers to continue analyzing the ABCD data and to help us improve the study by identifying potential problems like Bissett and colleagues have done. In light of the concerns raised by Bissett et al., some changes to the ABCD SST will be made (Table 3). The merits and demerits of changes to the stimuli presentation issues are being discussed and will be decided in the very near future. Fortunately, we believe that the specific concerns have relatively minor impact on the quality of the performance and neuroimaging data if the recommendations followed here (which does include some new participant exclusions) are followed. Consequently, we believe that any proposed changes to the task will not adversely affect longitudinal comparisons between the baseline data and subsequent assessments. By carefully phasing in any proposed changes we will be able to assess this empirically. More generally, we encourage researchers to contact the ABCD team promptly should their analyses raise concerns with any element of the assessment battery. Doing so ensures that misunderstandings can be avoided and any errors speedily corrected.

- Issue 1 (Different go stimulus durations across trials):
 - TBD. An ABCD working group with external expertise will weigh minimizing context violation concerns against altering an ongoing longitudinal study.
 - Violators (Stop Fail RT > Go RT) are identified in future data releases.
- Issue 2 (no Go stimuli at 0 SSD):
 - Resolution of this matter will follow from the resolution of Issue 1.
 - Future data releases include the numbers of 0SSD trials per participant.
- Issue 3 (coding error):
 - Participants with this error are identified in future data releases.
 - This error in the task will be corrected.
- Issue 4 (short Stop Signal durations when SSD is long):
 - The task will be changed to have a fixed Stop Signal duration.
- Issue 5 (non-uniform conditional trial probabilities):
 - No changes.
- Issue 6 (trial coding inaccuracies):
 - The task will be changed to correct these trial mislabelings.
 - All data releases have already corrected these errors.
- Issue 7 (SSD starts too short):
 - No changes.
- Issue 8 (low Stop trial probability):
 - No changes.

Table 3: Recommended task and data sharing changes.

References

Band GPH, van der Molen MW, Logan GD, (2003) Horse-race model simulations of the stop-signal procedure. *Acta Psychol. (Amst.)* 112, 105–142.

Bissett PG, Hagen MP, Poldrack RA, (Under Review) A cautionary note on stop-signal data from the Adolescent Brain Cognitive Development [ABCD] study. Preprint at <https://www.biorxiv.org/content/10.1101/2020.05.08.084707v1>

Bissett PG, Logan GD. (2012). Post-stop-signal adjustments: inhibition improves subsequent inhibition. *J Exp Psychol Learn Mem Cogn.* 38(4):955-966.

Bissett PG, Poldrack RA, Logan GD, (Under Review). Severe and pervasive violations of independence in response inhibition tasks. Preprint at <https://psyarxiv.com/kpa65/>

Bourque J, Spechler PA, Potvin S, et al. (2017). Functional Neuroimaging Predictors of Self-Reported Psychotic Symptoms in Adolescents. *Am J Psychiatry.* 174(6):566-575.

Casey BJ, Cannonier T, Conley MI, Cohen AO, Barch DM, Heitzeg MM, Soules ME, Teslovich T, Dellarco DV, Garavan H, Orr CA, Wager TD, Banich MT, Speer NK, Sutherland MT, Riedel MC, Dick AS, Bjork JM, Thomas KM, Chaarani B, Mejia MH, Hagler DJ Jr., Daniela Cornejo M, Sicut CS, Harms MP, Dosenbach NUF, Rosenberg M, Earl E, Bartsch H, Watts R, Polimeni JR, Kuperman JM, Fair DA, Dale AM; ABCD Imaging Acquisition Workgroup. (2018). The ABCD Study: Functional Imaging Acquisition across 21 Sites. *Developmental Cognitive Neuroscience.* 32:43-54.

Eisenberg IW, Bissett PG, Zeynep Enkavi A, et al. (2019) Uncovering the structure of self-regulation through data-driven ontology discovery. *Nat Commun.* 2019;10(1):2319.

Gulberti A, Arndt PA, Colonius H. (2014) Stopping eyes and hands: evidence for non-independence of stop and go processes and for a separation of central and peripheral inhibition. *Front Hum Neurosci.* 8:61.

Harlé KM, Zhang S, Ma N, Yu AJ, Paulus MP. (2016) Reduced Neural Recruitment for Bayesian Adjustment of Inhibitory Control in Methamphetamine Dependence. *Biol Psychiatry Cogn Neurosci Neuroimaging.* 1(5):448-459.

Leotti LA, Wager TD. (2010). Motivational influences on response inhibition measures. *Journal of Experimental Psychology: Human Perception and Performance.* 36(2):430-47.

Logan GD, Cowan WB, (1984) On the Ability to Inhibit Thought and Action: A Theory of an Act of Control. *Psychol. Rev.* 91, 295–327.

Poldrack RA, Congdon E, Triplett W, et al. (2016) A phenome-wide examination of neural and cognitive function. *Sci Data.* 3:160110.

Schumann G, Loth E, Banaschewski T, et al. (2010) The IMAGEN study: reinforcement-related behaviour in normal brain function and psychopathology. *Mol Psychiatry.* 2010;15(12):1128-1139.

Spechler PA, Allgaier N, Chaarani B, et al. (2019b) The initiation of cannabis use in adolescence is predicted by sex-specific psychosocial and neurobiological features. *Eur J Neurosci.*;50(3):2346-2356.

Spechler PA, Chaarani B, Orr C, et al. (2019a) Neuroimaging Evidence for Right Orbitofrontal Cortex Differences in Adolescents With Emotional and Behavioral Dysregulation. *J Am Acad Child Adolesc Psychiatry.* 58(11):1092-1103.

Verbruggen F, Aron AR, Band GP, et al. (2019). A consensus guide to capturing the ability to inhibit actions and impulsive behaviors in the stop-signal task. *Elife.*;8:e46323.

Whelan R, Conrod PJ, Poline JB, et al. (2012) Adolescent impulsivity phenotypes characterized by distinct brain networks. *Nat Neurosci.* 15(6):920-925.

Whelan R, Watts R, Orr CA, et al. (2014) Neuropsychosocial profiles of current and future adolescent alcohol misusers. *Nature.* 512(7513):185-189.